



Stochastic Mirror Descent Dynamics and Their Convergence in Monotone Variational Inequalities

Panayotis Mertikopoulos¹ · Mathias Staudigl² 

Received: 8 October 2017 / Accepted: 5 July 2018 / Published online: 18 July 2018
© The Author(s) 2018

Abstract

We examine a class of stochastic mirror descent dynamics in the context of monotone variational inequalities (including Nash equilibrium and saddle-point problems). The dynamics under study are formulated as a stochastic differential equation, driven by a (single-valued) monotone operator and perturbed by a Brownian motion. The system's controllable parameters are two variable weight sequences, that, respectively, pre- and post-multiply the driver of the process. By carefully tuning these parameters, we obtain global convergence in the ergodic sense, and we estimate the average rate of convergence of the process. We also establish a large deviations principle, showing that individual trajectories exhibit exponential concentration around this average.

Keywords Mirror descent · Variational inequalities · Saddle-point problems · Stochastic differential equations

Mathematics Subject Classification 90C25 · 90C33 · 90C47

1 Introduction

Dynamical systems governed by monotone operators play an important role in the fields of optimization (convex programming), game theory (Nash equilibrium and generalized Nash equilibrium problems), fixed point theory, partial differential equations and many other fields of applied mathematics. In particular, the study of the relationship between continuous- and discrete-time models has given rise to a vigor-

✉ Mathias Staudigl
m.staudigl@maastrichtuniversity.nl

Panayotis Mertikopoulos
panayotis.mertikopoulos@imag.fr

¹ CNRS, Inria, LIG, Univ. Grenoble Alpes, 38000 Grenoble, France

² Department of Quantitative Economics, School of Business and Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands

ous literature at the interface of these fields—see, e.g., [1] for a recent overview and [2] for connections to accelerated methods.

The starting point of much of this literature is that an iterative algorithm can be seen as a discretization of a continuous dynamical system. Doing so sheds new light on the properties of the algorithm, it provides Lyapunov functions, which are useful for its asymptotic analysis, and often leads to new classes of algorithms altogether. A classical example of this arises in the study of (projected) gradient descent dynamics and its connection with Cauchy’s steepest descent algorithm—or, more generally, in the relation between the mirror descent (MD) class of algorithms [3] and dynamical systems derived from Bregman projections and Hessian Riemannian metrics [4–6].

2 Problem Formulation and Related Literature

Throughout this paper, \mathcal{X} denotes a compact and convex subset of an n -dimensional real space $\mathcal{V} \cong \mathbb{R}^n$ with norm $\|\cdot\|$. We will also write $\mathcal{Y} \equiv \mathcal{V}^*$ for the dual of \mathcal{V} , $\langle y, x \rangle$ for the canonical pairing between $y \in \mathcal{Y}^*$ and $x \in \mathcal{V}$, and $\|y\|_* := \sup\{\langle y, x \rangle : \|x\| \leq 1\}$ for the dual norm of y in \mathcal{Y}^* . We denote the relative interior of \mathcal{X} by $\text{ri}(\mathcal{X})$, and its boundary by $\text{bd}(\mathcal{X})$.

In this paper, we are interested in deriving dynamical system approaches to solve monotone variational inequalities (VIs). To define them, let $v : \mathcal{X} \rightarrow \mathcal{Y}$ be a Lipschitz continuous monotone map, i.e.,

$$(H1) \quad \|v(x) - v(x')\|_* \leq L\|x - x'\| \quad \text{and} \quad \langle v(x) - v(x'), x - x' \rangle \geq 0,$$

for some $L > 0$ and all $x, x' \in \mathcal{X}$.

Throughout this paper, we will be interested in solving the Minty VI:

$$\text{Find } x_* \in \mathcal{X} \text{ such that } \langle v(x), x - x_* \rangle \geq 0 \text{ for all } x \in \mathcal{X}. \tag{MVI}$$

Since v is assumed continuous and monotone, this VI problem is equivalent to the Stampacchia VI:¹

$$\text{Find } x_* \in \mathcal{X} \text{ such that } \langle v(x_*), x - x_* \rangle \geq 0 \text{ for all } x \in \mathcal{X}. \tag{SVI}$$

When we need to keep track of \mathcal{X} and v explicitly, we will refer to (MVI) and/or (SVI) as $\text{VI}(\mathcal{X}, v)$. The solution set of $\text{VI}(\mathcal{X}, v)$ will be denoted as \mathcal{X}_* ; by standard results, \mathcal{X}_* is convex, compact and nonempty [9]. Below, we present a selected sample of examples and applications of VI problems; for a more extensive discussion, see [9–11].

Example 2.1 (Convex optimization) Consider the problem

$$\min_{x \in \mathcal{X}} f(x), \tag{Opt}$$

¹ Equivalence is understood in the sense that $x_* \in \mathcal{X}$ solves (MVI), if and only if it solves (SVI). See Proposition 3.1 in [7], and Proposition 3 in [8], which treats also extensions to Vector Variational Inequalities.

where $f : \mathcal{X} \rightarrow \mathbb{R}$ is convex and continuously differentiable on \mathcal{X} . If x_* is a solution of (Opt), first-order optimality gives

$$\langle \nabla f(x_*), x - x_* \rangle \geq 0 \quad \text{for all } x \in \mathcal{X}.$$

Since f is convex, $v = \nabla f$ is monotone, so (Opt) is equivalent to $\text{VI}(\mathcal{X}, \nabla f)$ [12].

Example 2.2 (Saddle-point problems) Let $\mathcal{X}^1 \subseteq \mathbb{R}^{n_1}$ and $\mathcal{X}^2 \subseteq \mathbb{R}^{n_2}$ be compact and convex, and let $U : \mathcal{X}^1 \times \mathcal{X}^2 \rightarrow \mathbb{R}$ be a smooth convex-concave function (i.e., $U(x^1, x^2)$ is convex in x^1 and concave in x^2). Then, the associated saddle-point (or min-max) problem is to determine the value of U , defined here as

$$\text{val} = \min_{x^1 \in \mathcal{X}^1} \max_{x^2 \in \mathcal{X}^2} U(x^1, x^2) = \max_{x^2 \in \mathcal{X}^2} \min_{x^1 \in \mathcal{X}^1} U(x^1, x^2). \tag{Val}$$

Existence of val follows directly from von Neumann’s minimax theorem. Moreover, letting

$$v(x^1, x^2) := (\nabla_{x^1} U(x^1, x^2), -\nabla_{x^2} U(x^1, x^2)), \tag{1}$$

it is easy to check that v is monotone as a map from $\mathcal{X} := \mathcal{X}^1 \times \mathcal{X}^2$ to $\mathbb{R}^{n_1+n_2}$ (because U is convex in its first argument and concave in the second). Then, as in the case of (Opt), first-order optimality implies that the saddle-points of (Val) are precisely the solutions of $\text{VI}(\mathcal{X}, v)$ [13].

Example 2.3 (Convex games) One of the main motivations for this paper comes from determining Nash equilibria of games with convex cost functions. To state the problem, let $\mathcal{N} := \{1, \dots, N\}$ be a finite set of players and, for each $i \in \mathcal{N}$, let $\mathcal{X}^i \subseteq \mathbb{R}^{n_i}$ be a compact convex set of actions that can be taken by player i . Given an action profile $x = (x^1, \dots, x^N) \in \mathcal{X} := \prod_i \mathcal{X}^i$, the cost for each player is determined by an associated cost function $c^i : \mathcal{X} \rightarrow \mathbb{R}$. The unilateral minimization of this cost leads to the notion of Nash equilibrium, defined here as an action profile $x_* = (x_*^i)_{i \in \mathcal{N}}$ such that

$$c^i(x_*) \leq c^i(x^i; x_*^{-i}) \quad \text{for all } x^i \in \mathcal{X}^i, i \in \mathcal{N}. \tag{NE}$$

Of particular interest to us is the case, where each c^i is smooth and individually convex in x^i . In this case, the profile $v(x) = (v^i(x))_{i \in \mathcal{N}}$ of individual gradients $v^i(x) := \nabla_{x^i} c^i(x)$ forms a monotone map and, by first-order optimality, the Nash equilibrium problem (NE) boils down to solving $\text{VI}(\mathcal{X}, v)$ [9,14].

In the rest of this paper, we will consider two important special cases of operators $v : \mathcal{X} \rightarrow \mathcal{V}^*$, namely:

1. Strictly monotone problems, i.e., when

$$\langle v(x') - v(x), x' - x \rangle \geq 0 \quad \text{with equality if and only if } x = x'.$$

2. *Strongly monotone* problems, i.e., when

$$\langle v(x') - v(x), x' - x \rangle \geq \gamma \|x' - x\|^2 \quad \text{for some } \gamma > 0.$$

Clearly, strong monotonicity implies strict monotonicity, which, in turn, implies ordinary monotonicity. In the case of convex optimization, strict (respectively, strong) monotonicity corresponds to strict (respectively, strong) convexity of the problem's objective function. Under either refinement, (MVI) admits a unique solution, which will be referred to as “the” solution of (MVI).

2.1 Contributions

Building on the above, this paper is concerned with a *stochastic* dynamical system resulting from Nesterov's well-known “dual-averaging” mirror descent algorithm [13], perturbed by noise, and/or random disturbances. Heuristically, this algorithm aggregates descent steps in the problem's (unconstrained) dual space, and then “mirrors” the result back to the problem's feasible region to obtain a candidate solution at each iteration. This “mirror step” is performed as in the classical setting of [3,4], but the dual aggregate is further post-multiplied by a variable parameter (thus turning “dual aggregates” into “dual averages”). Thanks to this averaging, the resulting algorithm is particularly suited for problems, where only noisy information is available to the optimizer, rendering it particularly useful for machine learning and engineering applications [15], even when the stochastic environment is not stationary [16].²

In more detail, the dynamics under study are formulated as a stochastic differential equation (SDE) driven by a (single-valued) monotone operator and perturbed by an Itô martingale noise process. As in Nesterov's original method [13], the dynamics' controllable parameters are two variable weight sequences, that, respectively, pre- and post-multiply the drift of the process: The first acts as a “step-size” of sorts, whereas the second can be seen as an “inverse temperature” parameter (as in simulated annealing). By carefully tuning these parameters, we are then able to establish the following results: First, if the intensity of the noise process decays with time, the dynamics converge to the (deterministic) solution of the underlying VI (cf. Sect. 4.2). Second, in the spirit of the ergodic convergence analysis of [13], we establish that this convergence can be achieved at an $\mathcal{O}(1/\sqrt{t})$ rate on average (Sect. 4.3).³ Finally, in Sect. 4.4, we establish a large deviations principle showing that, as far as ergodic convergence is concerned, the above convergence rate holds with (exponentially) high probability, not only in the mean.

Conceptually, our work has close ties to the literature on dynamical systems that arise in the solution of VIs, see, e.g., [2,4,19–22], and references therein. More specifically, a preliminary version of the dynamics considered in this paper was recently

² In the deterministic case, similar results are to be contrasted with the use of Gauss–Seidel methods for solving Variational and (generalized) Nash Equilibrium problems; for a survey, see [17,18].

³ Interestingly, the corresponding rate for the deterministic (noise-free) dynamics is $\mathcal{O}(1/t)$, indicating a substantial drop from the deterministic to the stochastic regime. This drop is consistent with the black-box convergence rate of Mirror Descent in (stochastic) VIs [13] and is due to the second-order Itô correction that appears in the stochastic case.

studied in the context of convex programming and gradient-like flows in [23,24]. The ergodic part of our analysis here extends the results of [24] to saddle-point problems and monotone variational inequalities, while the use of two variable weight sequences allows us to obtain almost sure convergence results without needing to rely on a parallel-sampling mechanism for variance reduction as in [23].

2.2 Stochastic Mirror Descent Dynamics

Mirror descent is an iterative optimization algorithm combining first-order oracle steps with a “mirror step” generated by a projection-type mapping. For the origins of the method, see [3]. The key ingredient defining this mirror step is a generalization of the Euclidean distance known as a “distance-generating” function:

Definition 2.1 We say that $h : \mathcal{X} \rightarrow \mathbb{R}$ is a *distance-generating function* on \mathcal{X} , if

- (a) h is *continuous*.
- (b) h is *strongly convex*, i.e., there exists some $\alpha > 0$ such that

$$h(\lambda x + (1 - \lambda)x') \leq \lambda h(x) + (1 - \lambda)h(x') - \frac{\alpha}{2}\lambda(1 - \lambda)\|x - x'\|^2,$$

for all $x, x' \in \mathcal{X}$ and all $\lambda \in [0, 1]$.

Given a distance-generating function on \mathcal{X} , its convex conjugate is given by

$$h^*(y) := \max_{x \in \mathcal{X}} \{ \langle y, x \rangle - h(x) \}, \quad y \in \mathcal{Y},$$

and the induced *mirror map* is defined as

$$Q(y) := \operatorname{argmax}_{x \in \mathcal{X}} \{ \langle y, x \rangle - h(x) \}.$$

Thanks to the strong convexity of h , $Q(y)$ is well-defined and single-valued for all $y \in \mathcal{Y}$. In particular, as illustrated in the examples below, it plays a role similar to that of a projection mapping:

Example 2.4 (Euclidean distance) If $h(x) = \frac{1}{2}\|x\|_2^2$, then the induced mirror map is the standard Euclidean projector

$$Q(y) = \operatorname{argmax}_{x \in \mathcal{X}} \left\{ \sum_{j=1}^n y_j x_j - \frac{1}{2} \sum_{j=1}^n x_j^2 \right\} = \operatorname{argmin}_{x \in \mathcal{X}} \|x - y\|_2^2.$$

Example 2.5 (Gibbs–Shannon entropy) If $\mathcal{X} = \{x \in \mathbb{R}_+^n : \sum_{j=1}^n x_j = 1\}$ is the unit simplex in \mathbb{R}^n , then the (negative) Gibbs–Shannon entropy $h(x) = \sum_{j=1}^n x_j \log x_j$ gives rise to the so-called *logit choice* map

$$Q(y) = \frac{(\exp(y_j))_{j=1}^n}{\sum_{k=1}^n \exp(y_k)}.$$

Example 2.6 (Fermi–Dirac entropy) If $\mathcal{X} = [0, 1]^n$ is the unit cube in \mathbb{R}^n , then the (negative) Fermi–Dirac entropy $h(x) = \sum_{j=1}^n [x_j \log(x_j) + (1 - x_j) \log(1 - x_j)]$ induces the so-called *logistic map*

$$Q(y) = \left(\frac{\exp(y_j)}{1 + \exp(y_j)} \right)_{j=1}^n .$$

For future reference, some basic properties of mirror maps are collected below:

Proposition 2.1 *Let h be a distance-generating function on \mathcal{X} . Then, the induced mirror map $Q: \mathcal{Y} \rightarrow \mathcal{X}$ satisfies the following properties:*

(a) $x = Q(y)$ if and only if $y \in \partial h(x)$, where

$$\partial h(x) := \{p \in \mathcal{V}^* : h(y) \geq h(x) + \langle p, y - x \rangle \quad \forall y \in \mathcal{X}\},$$

is the subgradient of h at x . In particular, $\text{im } Q = \text{dom } \partial h = \{x \in \mathcal{X} : \partial h(x) \neq \emptyset\}$.

(b) h^* is continuously differentiable on \mathcal{Y} and $\nabla h^*(y) = Q(y)$ for all $y \in \mathcal{Y}$.

(c) $Q(\cdot)$ is $(1/\alpha)$ -Lipschitz continuous.

The properties reported above are fairly standard in convex analysis; for a proof, see, e.g., [12, Theorem 12.60(b)]. Of particular importance is the identity $\nabla h^* = Q$, which provides a quick way of calculating Q in “prox-friendly” geometries (such as the examples discussed above).

Now, as mentioned above, mirror descent exploits the flexibility provided by a (not necessarily Euclidean) mirror map by using it to generate first-order steps along v . For concreteness, we will focus on the so-called “dual averaging” variant of mirror descent [13], defined here via the recursion

$$y_{t+1} = y_t - \lambda_t v(x_t), \quad x_{t+1} = Q(\eta_{t+1} y_{t+1}), \tag{2}$$

where:

- (1) $t = 0, 1, \dots$ denotes the stage of the process.
- (2) y_t is an auxiliary dual variable, aggregating first-order steps along v .⁴
- (3) λ_t is a variable step-size parameter, *pre-multiplying* the input at each stage.
- (4) η_t is a variable weight parameter, *post-multiplying* the dual aggregate y_t .⁵

Passing to continuous time, we obtain the *mirror descent dynamics*

$$dy(t) = -\lambda(t) v(x(t)) dt, \quad x(t) = Q(\eta(t)y(t)), \tag{MD}$$

⁴ The usual initialization is $y_0 = 0, x_0 = Q(0) = \underset{x \in \mathcal{X}}{\text{argmin}} h(x)$, but other initializations are possible.

⁵ The name “dual averaging” alludes to the choice $\lambda_t = 1, \eta_t = 1/t$: under this choice of parameters, x_t is a mirror projection of the “dual average” $y_t = t^{-1} \sum_{s=0}^{t-1} v(x_s)$.

with $\eta(t)$ and $\lambda(t)$ serving the same role as before (but now defined over all $t \geq 0$). In particular, our standing assumption for the parameters λ and η of (MD) will be that

$$(H2) \quad \eta(t) \text{ and } \lambda(t) \text{ are positive, } C^1\text{-smooth and nonincreasing.}$$

Heuristically, the assumptions above guarantee that the dual process $y(t)$ does not grow too large too fast, so blow-ups in finite time are not possible. Together with the basic convergence properties of the dynamics (MD), this is discussed in more detail in Sect. 3 below.

The primary case of interest in our paper is when the oracle information for $v(x)$ in (MD) is subject to noise, measurement errors and/or other stochastic disturbances. To account for such perturbations, we will instead focus on the *stochastic mirror descent dynamics*

$$dY(t) = -\lambda(t) [v(X(t)) dt + dM(t)], \quad X(t) = Q(\eta(t)Y(t)) \quad (\text{SMD})$$

where $M(t)$ is a continuous martingale with respect to some underlying stochastic basis $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$.⁶ In more detail, we assume for concreteness that the stochastic disturbance term $M(t)$ is an Itô process of the form

$$dM(t) = \sigma(X(t), t) \cdot dW(t),$$

where $W(t)$ is a d -dimensional Wiener process adapted to \mathcal{F}_t , and $\sigma(x, t)$ is an $n \times d$ matrix capturing the *volatility* of the noise process. Heuristically, the volatility matrix of $M(t)$ captures the intensity of the noise process, and the possible correlations between its components.

In terms of regularity, we will be assuming throughout that $\sigma(x, t)$ is measurable in t , as well as bounded, and Lipschitz continuous in x . Formally, we posit that there exists a constant $\ell > 0$ such that

$$(H3) \quad \sup_{x,t} \|\sigma(x, t)\| < \infty \quad \text{and} \quad \|\sigma(x', t) - \sigma(x, t)\| \leq \ell \|x' - x\|,$$

where

$$\|\sigma\| := \sqrt{\text{tr}[\sigma\sigma^\top]} = \sqrt{\sum_{i=1}^n \sum_{j=1}^d |\sigma_{ij}|^2}$$

denotes the Frobenius (matrix) norm of σ . In particular, (H3) implies that there exists a positive constant $\sigma_* \geq 0$ such that

$$\|\sigma(x, t)\|^2 \leq \sigma_*^2 \quad \text{for all } x \in \mathcal{X}, t \geq 0.$$

In what follows, it will be convenient to measure the intensity of the noise affecting (SMD) via σ_* ; of course, when $\sigma_* = 0$, we recover the noiseless, deterministic dynamics (MD).

⁶ We tacitly assume here that the filtration $(\mathcal{F}_t)_{t \geq 0}$ satisfies the usual conditions of right continuity and completeness, and carries a standard d -dimensional Wiener process $(W(t))_{t \geq 0}$.

3 Deterministic Analysis

To establish a reference standard, we first focus on the deterministic regime of (MD), i.e., when $M(t) \equiv 0$ in (SMD).

3.1 Global Existence

We begin with a basic well-posedness result of (MD).

Proposition 3.1 *Under Hypotheses (H1) and (H2), the dynamical system (MD) admits a unique solution from every initial condition $(s, y) \in \mathbb{R}_+ \times \mathcal{Y}$.*

Proof Let $A(t, y) := -\lambda(t)v(Q(\eta(t)y))$ for all $t \in \mathbb{R}_+, y \in \mathcal{Y}$. Clearly, $A(t, y)$ is jointly continuous in t and y . Moreover, by (H2), $\lambda(t)$ has bounded first derivative and $\eta(t)$ is nonincreasing, so both $\lambda(t)$ and $\eta(t)$ are Lipschitz continuous. Finally, by (H1), v is L -Lipschitz continuous, implying in turn that

$$\|A(t, y_1) - A(t, y_2)\|_* \leq \frac{L\eta(t)\lambda(t)}{\alpha} \|y_1 - y_2\|_* \quad \text{for all } y_1, y_2 \in \mathcal{Y},$$

where α is the strong convexity constant of h , and we used Proposition 2.1 to estimate the Lipschitz constant of Q . This shows that $A(t, y)$ is Lipschitz in y for all t , so existence and uniqueness of local solutions follows from the Picard–Lindelöf theorem. Hypothesis (H2) further guarantees that the Lipschitz constant of $A(t, \cdot)$ can be chosen uniformly in t , so these solutions can be extended for all $t \geq 0$. \square

Let $\mathbb{T} := \{(t, s) | 0 \leq s \leq t \leq \infty\}$. Based on the above, we may define a *nonautonomous semiflow* $Y : \mathbb{T} \times \mathcal{Y} \rightarrow \mathcal{Y}$ satisfying (i) $Y(s, s, y) = y$ for all $s \geq 0$, (ii) $\frac{\partial Y(t,s,y)}{\partial t} = A(t, Y(t, s, y))$ for all $(t, s, y) \in \mathbb{T} \times \mathcal{Y}$, and (iii) $Y(t, s, Y(s, r, y)) = Y(t, r, y)$ for $t \geq s \geq r \geq 0$. Since the dynamics will usually be started from an initial condition $(0, y) \in \mathbb{R}_+ \times \mathcal{Y}$, we will simplify the notation by writing $\phi(t, y) = Y(t, 0, y)$ for all $(t, y) \in \mathbb{R}_+ \times \mathcal{Y}$. The resulting trajectory in the primal space is denoted by $\xi(t, y) = Q(\eta(t)\phi(t, y))$. Note that, if $\lambda(t)$ and $\eta(t)$ are constant functions, then the mapping $\phi(t, y)$ is the (autonomous) *semiflow* of the dynamics (MD).

3.2 Convergence Properties and Performance

Now, to analyze the convergence of (MD), we will consider two “gap functions” quantifying the distance between the primal trajectory, and the solution set of (MVI):

- In the general case, we will focus on the *dual gap* function [25]:

$$g(x) := \max_{x' \in \mathcal{X}} \langle v(x'), x - x' \rangle.$$

By (H1) and the compactness of \mathcal{X} , it follows that $g(x)$ is continuous, nonnegative and convex; moreover, we have $g(x) = 0$ if and only if x is a solution of $\text{VI}(\mathcal{X}, v)$ [7, Proposition 3.1].

– For the saddle-point problem Ex. 2.2, we instead look at the *Nikaido–Isoda gap function* [26]:

$$G(p^1, p^2) := \max_{x^2 \in \mathcal{X}^2} U(p^1, x^2) - \min_{x^1 \in \mathcal{X}^1} U(x^1, p^2). \tag{3}$$

Since U is convex-concave, it is immediate that $G(p^1, p^2) \geq g(p^1, p^2)$, where the operator involved in the definition of the dual gap function is given by the saddle-point operator (1). However, it is still true that $G(p^1, p^2) = 0$ if and only if the pair (p^1, p^2) is a saddle-point. Since both gap functions vanish only at solutions of (MVI), we will prove trajectory convergence by monitoring the decrease of the relevant gap over time. This is achieved by introducing the so-called *Fenchel coupling* [14], an auxiliary energy function, defined as

$$F(x, y) := h(x) + h^*(y) - \langle y, x \rangle \quad \text{for all } x \in \mathcal{X}, y \in \mathcal{Y}.$$

Some key properties of F are summarized in the following proposition:

Proposition 3.2 ([14]) *Let h be a distance-generating function on \mathcal{X} . Then:*

- (a) $F(x, y) \geq \frac{\alpha}{2} \|Q(y) - x\|^2$ for all $x \in \mathcal{X}, y \in \mathcal{Y}$.
- (b) Viewed as a function of y , $F(x, y)$ is convex, differentiable, and its gradient is given by

$$\nabla_y F(x, y) = Q(y) - x.$$

- c) For all $x \in \mathcal{X}$ and all $y, y' \in \mathcal{Y}$, we have

$$F(x, y') \leq F(x, y) + \langle y' - y, Q(y) - x \rangle + \frac{1}{2\alpha} \|y' - y\|_*^2.$$

In the sequel, if there is no danger of confusion, we will use the more concise notation $x(t) = \xi(t, y)$ and $y(t) = \phi(t, y)$, for the unique solution to (MD) with initial condition $(0, y) \in \mathbb{R}_+ \times \mathcal{Y}$. Consider the averaged trajectory

$$\bar{x}(t) := \frac{\int_0^t \lambda(s)x(s) \, ds}{\int_0^t \lambda(s) \, ds} = \frac{1}{S(t)} \int_0^t \lambda(s)x(s) \, ds, \tag{4}$$

where $S(t) := \int_0^t \lambda(s) \, ds$. We then have the following convergence guarantee:

Proposition 3.3 *Suppose that (MD) is initialized at $(s, y) = (0, 0)$, with resulting trajectories $y(t) = \phi(t, 0)$ and $x(t) = \xi(t, 0)$. Then:*

$$g(\bar{x}(t)) \leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)S(t)}, \tag{5}$$

where $\bar{x}(t)$ is the averaged trajectory constructed in (4), and

$$\mathcal{D}(h; \mathcal{X}) := \max_{x, x' \in \mathcal{X}} \{h(x') - h(x)\} = \max h - \min h.$$

In particular, if (MVI) is associated with a convex-concave saddle-point problem as in Example 2.2, we have the guarantee:

$$G(\bar{x}(t)) \leq \frac{\mathcal{D}(h_1; \mathcal{X}^1) + \mathcal{D}(h_2; \mathcal{X}^2)}{\eta(t)S(t)}. \tag{6}$$

In both cases, whenever $\lim_{t \rightarrow \infty} \eta(t)S(t) = \infty$, $\bar{x}(t)$ converges to the solution set of VI(\mathcal{X}, v).

Proof Given some $p \in \mathcal{X}$, let $H_p(t) := \frac{1}{\eta(t)}F(p, \eta(t)y(t))$. Then, with Proposition 3.2, the fundamental theorem of calculus yields

$$H_p(t) - H_p(0) = - \int_0^t \lambda(s)\langle v(x(s)), x(s) - p \rangle ds - \int_0^t \frac{\dot{\eta}(s)}{\eta(s)^2} [h(p) - h(x(s))] ds,$$

and, after rearranging, we obtain

$$\begin{aligned} \int_0^t \lambda(s)\langle v(x(s)), x(s) - p \rangle ds &= H_p(0) - H_p(t) - \int_0^t \frac{\dot{\eta}(s)}{\eta(s)^2} [h(p) - h(x(s))] ds \\ &\leq H_p(0) + \mathcal{D}(h; \mathcal{X}) \left(\frac{1}{\eta(t)} - \frac{1}{\eta(0)} \right). \end{aligned} \tag{7}$$

Now, let $x_c := \operatorname{argmin}\{h(x) : x \in \mathcal{X}\}$ denote the ‘‘prox-center’’ of \mathcal{X} . Since $\eta(0) > 0$ and $y(0) = 0$ by assumption, we readily get

$$H_p(0) = \frac{F(p, 0)}{\eta(0)} = \frac{h(p) + h^*(0) - \langle 0, p \rangle}{\eta(0)} = \frac{h(p) - h(x_c)}{\eta(0)} \leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(0)}. \tag{8}$$

From the monotonicity of v , we further deduce that

$$g(\bar{x}(t)) \leq \frac{1}{S(t)} \max_{p \in \mathcal{X}} \int_0^t \lambda(s)\langle v(x(s)), x(s) - p \rangle ds. \tag{9}$$

Thus, substituting (8) in (7), maximizing over $p \in \mathcal{X}$ and plugging the result into (9) gives (5).

Suppose now that (MVI) is associated to a convex-concave saddle-point problem, as in Ex. 2.2. In this case, we can replicate the above analysis for each component $x^i(t)$, $i = 1, 2$, of $x(t)$ to obtain the basic bounds

$$\begin{aligned} \int_0^t \lambda(s)\langle \nabla_{x^1} U(x(s)), x^1(s) - p^1 \rangle ds &\leq \frac{\mathcal{D}(h_1; \mathcal{X}^1)}{\eta(t)}, \\ \int_0^t \lambda(s)\langle -\nabla_{x^2} U(x(s)), x^2(s) - p^2 \rangle ds &\leq \frac{\mathcal{D}(h_2; \mathcal{X}^2)}{\eta(t)}. \end{aligned}$$

Using the fact that U is convex-concave, this leads to the value-based bounds

$$\int_0^t \lambda(s)[U(x(s)) - U(p^1, x^2(s))] ds \leq \frac{\mathcal{D}(h_1; \mathcal{X}^1)}{\eta(t)},$$

$$\int_0^t \lambda(s)[U(x^1(s), p^2) - U(x(s))] ds \leq \frac{\mathcal{D}(h_2; \mathcal{X}^2)}{\eta(t)}.$$

Summing these inequalities, dividing by $S(t)$, and using Jensen’s inequality gives

$$U(\bar{x}^1(t), p^2) - U(p^1, \bar{x}^2(t)) \leq \frac{\mathcal{D}(h_1; \mathcal{X}^1) + \mathcal{D}(h_2; \mathcal{X}^2)}{\eta(t)S(t)}$$

The bound (6) then follows by taking the supremum over p^1 and p^2 , and using the definition of the Nikaido– Isoda gap function. □

The gap-based analysis of Proposition 3.3 can be refined further in the case of strongly monotone VIs.

Proposition 3.4 *Let x_* denote the (necessarily unique) solution of a γ -strongly monotone VI(\mathcal{X}, v). Then, with the same assumptions as in Proposition 3.3, we have*

$$\|\bar{x}(t) - x_*\|^2 \leq \frac{\mathcal{D}(h; \mathcal{X})}{\gamma} \frac{1}{\eta(t)S(t)}. \tag{10}$$

In particular, $\bar{x}(t)$ converges to x_ whenever $\lim_{t \rightarrow \infty} \eta(t)S(t) = \infty$.*

Proof By Jensen’s inequality, the strong monotonicity of v and the assumption that x_* solves VI(\mathcal{X}, v), we have:

$$\begin{aligned} \gamma \|\bar{x}(t) - x_*\|^2 &\leq \frac{\gamma}{S(t)} \int_0^t \lambda(s) \|x(s) - x_*\|^2 ds && \text{(Jensen)} \\ &\leq \frac{1}{S(t)} \int_0^t \lambda(s) \langle v(x(s)) - v(x_*), x(s) - x_* \rangle ds && (\gamma\text{-monotonicity}) \\ &\leq \frac{1}{S(t)} \int_0^t \lambda(s) \langle v(x(s)), x(s) - x_* \rangle ds && \text{(optimality of } x_*) \\ &\leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)S(t)}, \end{aligned}$$

where the last inequality follows as in the proof of Proposition 3.3. The bound (10) is then obtained by dividing both sides by γ . □

The two results above are in the spirit of classical ergodic convergence results for monotone VIs [13,27,28]. In particular, taking $\eta(t) = \sqrt{L/(2\alpha)}$ and $\lambda(t) = 1/(2\sqrt{t})$ gives the upper bound $g(\bar{x}(t)) \leq \mathcal{D}(h; \mathcal{X})\sqrt{L/(\alpha t)}$, which is of the same order as the $\mathcal{O}(1/\sqrt{t})$ guarantees obtained in the references above. However, the bound (9) does not have a term which is antagonistic to $\eta(t)$ or $\lambda(t)$, so, if (MD) is run with constant λ and η , we get an $\mathcal{O}(1/t)$ bound for $g(\bar{x}(t))$ (and/or $\|\bar{x}(t) - x_*\|$ in the case

of strongly monotone VIs).⁷ This suggests an important gap between continuous and discrete time; for a similar phenomenon in the context of online convex optimization, see the regret minimization analysis of [29].

We close this section with a (nonergodic) trajectory convergence result for strictly monotone problems. For any path $X(\cdot) : \mathbb{R}_+ \rightarrow \mathcal{X}$, call the *limit set*

$$\mathcal{L}\{X(\cdot)\} := \bigcap_{t \geq 0} \text{cl}[X([t, \infty))].$$

Proposition 3.5 *Let x_* denote the (necessarily unique) solution of a γ -strongly monotone VI(\mathcal{X}, v). Suppose that Hypotheses (H1) and (H2) hold, and the parameters λ and η of (MD) satisfy*

$$\inf_t \lambda(t) > 0 \quad \text{and} \quad \inf_t \eta(t) > 0.$$

Then, $\lim_{t \rightarrow \infty} \xi(t, y) = x_$, for any $y \in \mathcal{Y}$.*

Proof Let $x(t) := \xi(t, y)$ for $t \geq 0$, and assume that $\hat{x} \in \mathcal{L}\{x(\cdot)\}$, but $\hat{x} \neq x_*$. Then, by assumption, there exists an open neighborhood O of \hat{x} and a positive constant $a > 0$ such that

$$\langle v(x), x - x_* \rangle \geq a \quad \text{for all } x \in O.$$

Furthermore, since \hat{x} is an accumulation point of $x(t)$, there exists an increasing sequence $(t_k)_{k \in \mathbb{N}}$ such that $t_k \uparrow \infty$ and $x(t_k) \rightarrow \hat{x}$ as $k \rightarrow \infty$. Thus, relabeling indices if necessary, we may assume without loss of generality that $x(t_k) \in O$ for all $k \in \mathbb{N}$. Now, for all $\varepsilon > 0$, we have

$$\begin{aligned} \|x(t_k + \varepsilon) - x(t_k)\| &= \|Q(Y(t_k + \varepsilon)) - Q(Y(t_k))\| \\ &\leq \frac{1}{\alpha} \|Y(t_k + \varepsilon) - Y(t_k)\|_* \\ &\leq \frac{1}{\alpha} \int_{t_k}^{t_k + \varepsilon} \lambda(s) \|v(x(s))\|_* \, ds \\ &\leq \frac{1}{\alpha} \max_{x \in \mathcal{X}} \|v(x)\|_* \int_{t_k}^{t_k + \varepsilon} \lambda(s) \, ds \\ &\leq \frac{\varepsilon \bar{\lambda}}{\alpha} \max_{x \in \mathcal{X}} \|v(x)\|_*, \end{aligned}$$

where $\bar{\lambda} := \lambda(0)$ denotes the maximum value of $\lambda(t)$. As this bound does not depend on k , we can choose $\varepsilon > 0$ small enough so that $x(t_k + s) \in O$ for all $s \in [0, \varepsilon]$ and all $k \in \mathbb{N}$. Thus, letting $H(t) := \eta(t)^{-1} F(x_*, \eta(t)y(t))$, and using (7), we obtain

$$H(t_n) - H(t_0) = - \sum_{k=1}^n \int_{t_{k-1}}^{t_k} \lambda(s) \langle v(x(s)), x(s) - x_* \rangle \, ds + \mathcal{D}(h; \mathcal{X}) \left(\frac{1}{\eta(t_n)} - \frac{1}{\eta(t_0)} \right)$$

⁷ In fact, even faster convergence can be guaranteed if (MD) is run with increasing $\lambda(t)$. In that case however, well posedness is not immediately guaranteed, so we do not consider increasing λ here.

$$\begin{aligned} &\leq -a\underline{\lambda} \sum_{k=1}^n (t_k - t_{k-1}) + \mathcal{D}(h; \mathcal{X}) \left(\frac{1}{\eta(t_n)} - \frac{1}{\eta(t_0)} \right) \\ &= -a\varepsilon \underline{\lambda} n + \mathcal{D}(h; \mathcal{X}) \left(\frac{1}{\eta(t_n)} - \frac{1}{\eta(t_0)} \right), \end{aligned}$$

where we have set $\underline{\lambda} := \inf_t \lambda(t) > 0$. Given that $\inf_t \eta(t) > 0$, the above implies that $\lim_{n \rightarrow \infty} H(t_n) = -\infty$, contradicting the fact that $F(x_*, y) \geq 0$ for all $y \in \mathcal{Y}$. This implies that $\hat{x} = x_*$; by compactness, $\mathcal{L}\{x(\cdot)\} \neq \emptyset$, so our claim follows. \square

4 Analysis of the Stochastic Dynamics

4.1 Global Existence

In this section, we turn to the stochastic system (SMD). As in the noise-free analysis of the previous section, we begin with a well-posedness result, stated for simplicity for deterministic initial conditions.

Proposition 4.1 *Fix an initial condition $(s, y) \in \mathbb{R}_+ \times \mathcal{Y}$. Then, under Hypotheses (H1)–(H3), and up to a \mathbb{P} -null set, the stochastic dynamics (SMD) admit a unique strong solution $(Y(t))_{t \geq s}$ such that $Y(s) = y$.*

Proof Let $B(t, y) := -\lambda(t)\sigma(Q(\eta(t)y), t)$ so (SMD) can be written as

$$dY(t) = A(t, Y(t)) dt + B(t, Y(t)) dW(t), \quad (11)$$

with $A(t, y)$ defined as in the proof of Proposition 3.1. By (H2) and (H3), $B(t, y)$ inherits the boundedness and regularity properties of σ ; in particular, Hypotheses (H2) and (H3), together with Proposition 2.1c), imply that $B(t, y)$ is uniformly Lipschitz in y . Under Hypotheses (H1) and (H3), $A(t, y)$ is also uniformly Lipschitz in y (cf. the proof of Proposition 3.1). Our claim then follows by standard results in the well posedness of stochastic differential equations [30, Theorem 3.4]. \square

We denote by $Y(t, s, y)$ the unique strong solution of the Itô stochastic differential equation (11), with initial condition $(s, y) \in \mathbb{R}_+ \times \mathcal{Y}$. As in the deterministic case, we are mostly interested in the process starting from the initial condition $(0, y)$, in which case we abuse the notation by writing $Y(t, y) = Y(t, 0, y)$. The corresponding primal trajectories are generated by applying the mirror map Q to the dual trajectories, so $X(t, y) = Q(\eta(t)Y(t, y))$, for all $(t, y) \in \mathbb{R}_+ \times \mathcal{Y}$. If there is no danger of confusion, we will consistently suppress the dependence on the initial position $y \in \mathcal{Y}$ in both random processes. Clearly, if $\lambda(t)$ and $\eta(t)$ are constant function, the solutions of (SMD) are time-autonomous.

We now give a brief overview of the results we obtain in this section. First, in Sect. 4.2, we use the theory of *asymptotic pseudo-trajectories* (APTs), developed by Benaïm and Hirsch [31], to establish almost sure trajectory convergence of (SMD) to the solution of $\text{VI}(\mathcal{X}, v)$, provided that v is strictly monotone and the oracle noise in (SMD) is vanishing at a rather slow, logarithmic rate. This strong convergence result

relies heavily on the shadowing property of the dual trajectory, and its deterministic counterpart $\phi(t, y)$ (see Sect. 4.2). On the other hand, if the driving noise process is persistent, we cannot expect the primal trajectory $X(t)$ to converge—some averaging has to be done in this case. Thus, following a long tradition on ergodic convergence for mirror descent, we investigate in Sect. 4.3 the asymptotics of a weighted time-average of $X(t)$. Finally, we complement our ergodic convergence results with a large deviation principle showing that the ergodic average of $X(t)$ is exponentially concentrated around its mean (Sect. 4.4).

4.2 The Small Noise Limit

We begin with the case, where the oracle noise in (SMD) satisfies the asymptotic decay condition $\|\sigma(x, t)\| \leq \beta(t)$ for some nonincreasing function $\beta: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$(H4) \quad \int_0^\infty \exp\left(-\frac{c}{\beta^2(t)}\right) dt < \infty \quad \text{for all } c > 0.$$

For instance, this condition is trivially satisfied if $\sigma(x, t)$ vanishes at a logarithmic rate, i.e., $\beta(t) = o(1/\sqrt{\log(t)})$. For technical reasons, we will also need the additional ‘‘Fenchel reciprocity’’ condition

$$(H5) \quad F(p, y_n) \rightarrow 0 \quad \text{whenever } Q(y_n) \rightarrow p.$$

Under the decay rate requirement (H4), and working for simplicity with constant $\eta(t) = \lambda(t) = 1$, the results of [31, Proposition 4.1] imply that any strong solution $Y(t)$ of (SMD) is an (APT) of the deterministic dynamics (MD) in the following sense:

Definition 4.1 Assume that $\eta(t) = \lambda(t) = 1$, for all $t \geq 0$. Let $\phi: \mathbb{R}_+ \times \mathcal{Y} \rightarrow \mathcal{Y}$, $(t, y) \mapsto \phi(t, y)$, denote the semiflow induced by (MD) on \mathcal{Y} . A continuous curve $Y: \mathbb{R}_+ \rightarrow \mathcal{Y}$ is said to be an *asymptotic pseudo-trajectory* (APT) of (MD), if

$$\lim_{t \rightarrow \infty} \sup_{0 \leq s \leq T} \|Y(t + s) - \phi(s, Y(t))\|_* = 0 \quad \text{for all } T > 0. \quad (\text{APT})$$

In words, Definition 4.1 states that an APT of (MD) tracks the solutions of (MD) to arbitrary accuracy over arbitrarily long time windows. Thanks to this property, we are able to establish the following global convergence theorem for (SMD) with vanishing oracle noise:

Theorem 4.1 Assume that v is strictly monotone, and let x_* denote the (necessarily unique) solution of $\text{VI}(\mathcal{X}, v)$. If Hypotheses (H1)–(H5) hold, and (SMD) is run with $\lambda(t) = \eta(t) = 1$, we have

$$\mathbb{P}\left(\lim_{t \rightarrow \infty} \|X(t, y) - x_*\| = 0\right) = 1 \quad \forall y \in \mathcal{Y}.$$

The proof of Theorem 4.1 requires some auxiliary results, which we provide below. We begin with a strong recurrence result for neighborhoods of the (unique) solution x_* of VI(\mathcal{X}, v) under (MD):

Lemma 4.1 *With assumptions as in Theorem 4.1, let \mathcal{O} be an open neighborhood of x_* in \mathcal{X} and let $\xi(t, y) = Q(\eta(t)\phi(t, y))$. Define the stopping time*

$$t_{\mathcal{O}}(y) := \inf\{t \geq 0 : \xi(t, y) \in \mathcal{O}\}.$$

Then, $t_{\mathcal{O}}(y) < \infty$ for all $y \in \mathcal{Y}$.

Proof Fix the initialization $y \in \mathcal{Y}$ of (MD), and let $y(t) := \phi(t, y)$, $x(t) := Q(\phi(t, y))$ denote the induced solutions of (MD), and let $H(t) := F(x_*, y(t))$. Then, by Proposition 3.2, and the chain rule applied to (MD), we get

$$H(t) = H(0) - \int_0^t \langle v(x(s)), x(s) - x_* \rangle ds.$$

Since v is strictly monotone and x_* solves VI(\mathcal{X}, v), there exists some $a \equiv a_{\mathcal{O}} > 0$ such that

$$\langle v(x), x - x_* \rangle \geq a \quad \text{for all } x \in \mathcal{X} \setminus \mathcal{O}.$$

Hence, if $t_{\mathcal{O}}(y) = \infty$, we would have

$$H(t) \leq H(0) - at \quad \text{for all } t \geq 0,$$

implying in turn that $\lim_{t \rightarrow \infty} H(t) = -\infty$. This contradicts the fact that $H(t) \geq 0$, so we conclude that $t_{\mathcal{O}}(y) < \infty$. □

Next, we extend this result to the stochastic regime:

Lemma 4.2 *With assumptions as in Theorem 4.1, let \mathcal{O} be an open neighborhood of x_* in \mathcal{X} and define the stopping time*

$$\tau_{\mathcal{O}}(y) := \inf\{t \geq 0 : X(t, y) \in \mathcal{O}\},$$

Then, $\tau_{\mathcal{O}}(y)$ is almost surely finite for all $y \in \mathcal{Y}$.

Proof Suppose there exists some initial condition $y_0 \in \mathcal{Y}$, such that $\mathbb{P}(\tau_{\mathcal{O}}(y_0) = \infty) > 0$. Then, there exists a measurable set $\Omega_0 \subseteq \Omega$, with $\mathbb{P}(\Omega_0) > 0$, and such that $\tau_{\mathcal{O}}(\omega, y_0) = \infty$ for all $\omega \in \Omega_0$. Now, define $H(t) := F(x_*, Y(t, y_0))$ and set $X(t) = X(t, y_0)$. By the weak Itô lemma (33) proven in Sect. 5, we get

$$H(t) - H(0) \leq - \int_0^t \langle v(X(s)), X(s) - x_* \rangle ds + \frac{1}{2\alpha} \int_0^t \|\sigma(X(s), s)\|^2 ds + I_{x_*}(t)$$

where $I_{x_*}(t) := \int_0^t \langle X(s) - x_*, \sigma(X(s)) \cdot dW(s) \rangle$ is a continuous local martingale. Since v is strictly monotone, the same reasoning as in the proof of Lemma 4.1 yields

$$H(t) \leq H(0) - at + I_{x_*}(t) + \frac{\sigma_*^2}{2\alpha}$$

for some $a \equiv a_{\mathcal{O}} > 0$ and for all $t \in [0, \tau_{\mathcal{O}}(y))$. Furthermore, by an argument based on the law of the iterated logarithm and the Dambis–Dubins–Schwarz time-change theorem for martingales as in the proof of Theorem 4.2, we get

$$I_{x_*}(t)/t \rightarrow 0 \text{ almost surely as } t \rightarrow \infty.$$

Combining this with the estimate for $H(t)$ above, we get $\lim_{t \rightarrow \infty} H(t) = -\infty$ for \mathbb{P} -almost all $\omega \in \Omega_0$. This contradicts the fact that $H(t) \geq 0$, and our claim follows. \square

The above result shows that the primal process $X(t)$ hits any neighborhood of x_* in finite time (a.s.). Thanks to this important recurrence property, we are finally in a position to prove Theorem 4.1:

Proof of Theorem 4.1 Fix some $\varepsilon > 0$, and let $N_\varepsilon := \{x = Q(y) : F(x_*, y) < \varepsilon\}$. Let $y \in \mathcal{Y}$ be arbitrary. We first claim, that there exists a deterministic time $T \equiv T(\varepsilon)$ such that $F(x_*, \phi(T, y)) \leq \max\{\varepsilon, F(x_*, y) + \varepsilon\}$. Indeed, consider the hitting time

$$t_\varepsilon(y) := \inf\{t \geq 0 : x(t) \in N_\varepsilon\},$$

where $x(t) := Q(\phi(t, y))$. By Hypothesis (H5), N_ε contains a neighborhood of x_* ; hence, by Lemma 4.1, we have $t_\varepsilon(y) < \infty$. Moreover, observe that

$$\frac{d}{dt} F(x_*, \phi(t, y)) = -\langle v(x(t)), x(t) - x_* \rangle \leq 0 \text{ for all } y \in \mathcal{Y}. \tag{12}$$

The strict monotonicity of v and the fact that x_* solves (MVI) imply that there exists a positive constant $\kappa \equiv \kappa_\varepsilon > 0$ such that $\langle v(x), x - x_* \rangle \geq \kappa$ for all $x \in \mathcal{X} \setminus N_\varepsilon$. Hence, combining this with (12), we readily see that

$$F(x_*, \phi(t, y)) - F(x_*, y) \leq -\kappa t \text{ for all } t \in [0, t_\varepsilon(y)).$$

Now, set $T = \varepsilon/\kappa$. If $T < t_\varepsilon(y)$, we immediately conclude that

$$F(x_*, \phi(T, y)) - F(x_*, y) \leq -\varepsilon.$$

Otherwise, if $T \geq t_\varepsilon(y)$, we again use the descent property (12) to get

$$F(x_*, \phi(T, y)) \leq F(x_*, \phi(t_\varepsilon(y), y)) \leq \varepsilon.$$

In both cases we have $F(x_*, \phi(T, y)) \leq \max\{\varepsilon, F(x_*, y) - \varepsilon\}$, as claimed.

To proceed, pick $\delta \equiv \delta_\varepsilon > 0$ such that

$$\delta_\varepsilon \operatorname{diam}(\mathcal{X}) + \frac{\delta_\varepsilon^2}{2\alpha} < \varepsilon, \tag{13}$$

where $\operatorname{diam}(\mathcal{X}) := \max\{\|x' - x\|_2 : x, x' \in \mathcal{X}\}$ denotes the Euclidean diameter of \mathcal{X} . By Proposition 4.1 of [31], the strong solution Y of (11) (viewed as a stochastic flow) is an APT of the deterministic semiflow ϕ with probability 1. Hence, we can choose an (a.s.) finite random time θ_ε such that $\sup_{s \in [0, T]} \|Y(t + s) - \phi(s, Y(t))\|_* \leq \delta_\varepsilon$ for all $t \geq \theta_\varepsilon$. Combining this with item (c) of Proposition 3.2, we then get

$$\begin{aligned} F(x_*, Y(t + s, y)) &\leq F(x_*, \phi(s, Y(t, y))) \\ &\quad + \langle Y(t + s, y) - \phi(s, Y(t, y)), Q(\phi(s, Y(t, y))) - x_* \rangle \\ &\quad + \frac{1}{2\alpha} \|Y(t + s, y) - \phi(s, Y(t, y))\|_*^2 \\ &\leq F(x_*, \phi(s, Y(t, y))) + \delta_\varepsilon \operatorname{diam}(\mathcal{X}) + \frac{\delta_\varepsilon^2}{2\alpha} \\ &\leq F(x_*, \phi(s, Y(t, y))) + \varepsilon, \end{aligned}$$

where the last inequality follows from the estimate (13).

Now, choose a random time $T_0 \geq \max\{\theta_\varepsilon(y), t_\varepsilon(y)\}$ and $T = \varepsilon/\kappa$ as above. Then, by definition, we have $F(x_*, Y(T_0, y)) \leq 2\varepsilon$ with probability 1. Hence, for all $s \in [0, T]$, we get

$$F(x_*, Y(T_0 + s, y)) \leq F(x_*, \phi(s, Y(T_0, y))) + \varepsilon \leq F(x_*, Y(T_0, y)) + \varepsilon \leq 3\varepsilon.$$

Since $F(x_*, \phi(T, Y(T_0, y))) \leq \max\{\varepsilon, F(x_*, Y(T_0, y)) - \varepsilon\} \leq \varepsilon$, we also get

$$\begin{aligned} F(x_*, Y(T_0 + T + s, y)) &\leq F(x_*, \phi(s, Y(T_0 + T, y))) + \varepsilon \\ &\leq F(x_*, Y(T_0 + T, y)) + \varepsilon \\ &\leq 3\varepsilon, \end{aligned}$$

and hence

$$F(x_*, Y(T_0 + s, y)) \leq 3\varepsilon \quad \text{for all } s \in [T, 2T].$$

Using this as the basis for an induction argument, we readily get

$$F(x_*, Y(T_0 + s, y)) \leq 3\varepsilon \quad \text{for all } s \in [nT, (n + 1)T],$$

with probability 1. Since ε was arbitrary, we obtain $F(x_*, Y(t, y)) \rightarrow 0$, implying in turn that $X(t) \rightarrow x_*$ (a.s.) by Proposition 3.2. □

4.3 Ergodic Convergence

We now proceed with an ergodic convergence result, in the spirit of Proposition 3.3. The results presented in this section are derived under the assumption that (SMD) is started with the initial conditions $(s, y) = (0, 0)$. This is only done to make the presentation clearer; see Remark 4.1.

Set $S(t) := \int_0^t \lambda(s) \, ds$, $L(t) := \sqrt{\int_0^t \lambda^2(s) \, ds}$, and let

$$\bar{X}(t) := \frac{1}{S(t)} \int_0^t \lambda(s) X(s) \, ds,$$

denote the “ergodic average” of $X(t) = X(t, 0, 0)$.

Theorem 4.2 *Under Hypotheses (H1)–(H3), we have:*

$$g(\bar{X}(t)) = \mathcal{O}\left(\frac{1}{\eta(t)S(t)}\right) + \mathcal{O}\left(\frac{\int_0^t \lambda^2(s)\eta(s) \, ds}{S(t)}\right) + \mathcal{O}\left(\frac{L(t)\sqrt{\log \log L(t)}}{S(t)}\right), \tag{14}$$

with probability 1. In particular, $\bar{X}(t)$ converges (a.s.) to the solution set of VI(\mathcal{X}, v) provided that a) $\lim_{t \rightarrow \infty} \eta(t)S(t) = \infty$; and b) $\lim_{t \rightarrow \infty} \eta(t)\lambda(t) = 0$.

Before discussing the proof of Theorem 4.2, it is worth noting the interplay between the two variable weight parameters, $\lambda(t)$ and $\eta(t)$. In particular, if (SMD) is run with weight sequences of the form $1/t^q$ for some $q > 0$, we obtain:

Corollary 4.1 *Suppose that (SMD) is run with $\lambda(t) = (1 + t)^{-a}$ and $\eta(t) = (1 + t)^{-b}$ for some $a, b \in [0, 1]$. Then, with assumptions as in Theorem 4.2, we have:*

$$g(\bar{X}(t)) = \tilde{\mathcal{O}}\left(t^{a+b-1}\right) + \tilde{\mathcal{O}}\left(t^{-a-b}\right) + \tilde{\mathcal{O}}\left(t^{-1/2}\right). \tag{15}$$

In the above, the $\tilde{\mathcal{O}}(\cdot)$ notation signifies “ $\mathcal{O}(\cdot)$ up to logarithmic factors”.⁸ Up to such factors, (15) is optimized when $a + b = 1/2$; if these factors are to be considered, any choice with $a + b = 1/2$ and $b > 0$ gives the same rate of convergence, indicating that the role of the post-multiplication factor $\eta(t)$ is crucial to finetune the convergence rate of (SMD). We find this observation particularly appealing, as it is reminiscent of Nesterov’s remark that “running the discrete-time algorithm (2) with the *best* step-size strategy λ_t and fixed η [...] gives the same (infinite) constant as the corresponding strategy for fixed λ and variable η_t ” [13, p. 224].

The proof of Theorem 4.2 relies crucially on the following lemma, which provides an explicit estimate for the decay rate of the employed gap functions.

Lemma 4.3 *If (SMD) is initialized at $(0, 0)$, and Hypotheses (H1)–(H3) hold, then:*

$$g(\bar{X}(t)) \leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)S(t)} + \frac{\sigma_*^2 \int_0^t \lambda^2(s)\eta(s) \, ds}{2\alpha S(t)} + \frac{I(t)}{S(t)} \tag{16}$$

⁸ More precisely, we write $f(x) = \tilde{\mathcal{O}}(g(x))$ when $f(x) = \mathcal{O}(g(x) \log^k g(x))$ for some $k > 0$.

where $I(t) := \sup_{p \in \mathcal{X}} I_p(t)$ and

$$I_p(t) := \int_0^t \lambda(s) \langle p - X(s), \sigma(X(s), s) \cdot dW(s) \rangle. \tag{17}$$

If (MVI) is associated with a convex-concave saddle-point problem as in Example 2.2, we have

$$G(\bar{X}^1(t), \bar{X}^2(t)) \leq \frac{\mathcal{D}_{sp}}{\eta(t)S(t)} + \frac{\sigma_*^2}{2\alpha_{sp}} \frac{\int_0^t \lambda^2(s)\eta(s) ds}{S(t)} + \frac{J(t)}{S(t)},$$

where we have set $\mathcal{D}_{sp} := \mathcal{D}(h_1; \mathcal{X}^1) + \mathcal{D}(h_2; \mathcal{X}^2)$, $1/\alpha_{sp} := 1/\alpha_1 + 1/\alpha_2$, and $J(t) := \sup_{p^1 \in \mathcal{X}^1, p^2 \in \mathcal{X}^2} \{I_{p^1}(t) + I_{p^2}(t)\}$.

Remark 4.1 The initialization assumption in Lemma 4.3 is not crucial: We only make it to simplify the explicit expression (16). If (SMD) is initialized at a different point, the proof of Lemma 4.3 shows that the bound (16) is correct only up to $\mathcal{O}(1/S(t))$. Since all terms in (16) are no faster than $\mathcal{O}(1/S(t))$, initialization plays no role in the proof of Theorem 4.2 below.

Proof of Lemma 4.3 Fix some $p \in \mathcal{X}$, and let $H_p(t) := \eta(t)^{-1} F(p, \eta(t)Y(t))$, as in the proof of Proposition 3.3. Then, by the weak Itô formula (33) in Sect. 5, we have

$$\begin{aligned} H_p(t) &\leq H_p(0) - \int_0^t \frac{\dot{\eta}(s)}{\eta(s)^2} H_p(s) ds + \frac{1}{\eta(t)} \int_0^t \langle X(s) - p, \dot{\eta}(s)Y(s) \rangle ds \\ &\quad + \int_0^t \langle X(s) - p, dY(s) \rangle + \frac{1}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \|\sigma(X(s), s)\|^2 ds. \end{aligned} \tag{18}$$

To proceed, let

$$R_p(t) := \int_0^t \lambda(s) \langle v(X(s)), X(s) - p \rangle ds, \tag{19}$$

so

$$\begin{aligned} \int_0^t \langle X(s) - p, dY(s) \rangle &= - \int_0^t \lambda(s) \langle X(s) - p, v(X(s)) \rangle ds + dM(s) \\ &= -R_p(t) + I_p(t), \end{aligned}$$

with $I_p(t)$ given by (17). Then, rearranging and bounding the second term of (18) as in the proof of Proposition 3.3, we obtain

$$\begin{aligned} R_p(t) &\leq H_p(0) - H_p(t) + \mathcal{D}(h, \mathcal{X}) \left(\frac{1}{\eta(t)} - \frac{1}{\eta(0)} \right) \\ &\quad + I_p(t) + \frac{1}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \|\sigma(X(s), s)\|^2 ds \\ &\leq H_p(0) + \mathcal{D}(h; \mathcal{X}) \left(\frac{1}{\eta(t)} - \frac{1}{\eta(0)} \right) + I_p(t) + \frac{\sigma_*^2}{2\alpha} \int_0^t \lambda^2(s)\eta(s) ds. \end{aligned}$$

With (SMD) initialized at $y = 0$, Eq. (8) gives $H_p(0) \leq \mathcal{D}(h; \mathcal{X})/\eta(0)$. Thus, by Jensen’s inequality and the monotonicity of v , we get

$$\begin{aligned} \langle v(p), \bar{X}(t) - p \rangle &= \frac{1}{S(t)} \int_0^t \lambda(s) \langle v(p), X(s) - p \rangle \, ds \\ &\leq \frac{1}{S(t)} \int_0^t \lambda(s) \langle v(X(s)), X(s) - p \rangle \, ds = \frac{R_p(t)}{S(t)} \\ &\leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)S(t)} + \frac{\sigma_*^2 \int_0^t \lambda^2(s)\eta(s) \, ds}{2\alpha S(t)} + \frac{I_p(t)}{S(t)}. \end{aligned} \tag{20}$$

The bound (16) then follows by noting that $g(\bar{X}(t)) = \max_{p \in \mathcal{X}} \langle v(p), \bar{X}(t) - p \rangle$.

Now, assume that (MVI) is associated to a convex-concave saddle-point problem as in Ex. 2.2. As in the proof of Proposition 3.3, we first replicate the analysis above for each component of the problem, and we then sum the two components to get an overall bound for the Nikaido–Isoda gap function G . Specifically, applying (20) to (1), we readily get

$$\int_0^t \lambda(s) \langle v^i(X(s)), X^i(s) - p^i \rangle \leq \frac{\mathcal{D}(h_i; \mathcal{X}^i)}{\eta(t)S(t)} + \frac{\sigma_*^2 \int_0^t \lambda^2(s)\eta(s) \, ds}{2\alpha^i S(t)} + \frac{I_{p^i}(t)}{S(t)}, \tag{21}$$

where $i \in \{1, 2\}$. Moreover, Jensen’s inequality yields

$$\begin{aligned} U(\bar{X}^1(t), p^2) - U(p^1, \bar{X}^2(t)) &\leq \frac{1}{S(t)} \int_0^t \lambda(s) \left[U(X^1(s), p^2) - U(p^1, X^2(s)) \right] \, ds \\ &\leq \frac{1}{S(t)} \int_0^t \lambda(s) \langle \nabla_{x^1} U(X(s)), X^1(s) - p^1 \rangle \, ds \\ &\quad - \frac{1}{S(t)} \int_0^t \lambda(s) \langle \nabla_{x^2} U(X(s)), X^2(s) - p^2 \rangle \, ds \\ &\leq \frac{\mathcal{D}_{sp}}{\eta(t)S(t)} + \frac{\sigma_*^2 \int_0^t \lambda^2(s)\eta(s) \, ds}{2\alpha_{sp} S(t)} + \frac{I_{p^1}(t) + I_{p^2}(t)}{S(t)}, \end{aligned}$$

with the last inequality following from (21). Our claim then follows by maximizing over (p^1, p^2) and recalling the definition (3) of the Nikaido–Isoda gap function. \square

Clearly, the crucial unknown in the bound (16) is the stochastic term $I(t)$. To obtain convergence of $\bar{X}(t)$ to the solution set of $\text{VI}(\mathcal{X}, v)$, the term $I(t)$ must grow slower than $S(t)$. As we show now, this is indeed the case:

Proof of Theorem 4.2 By Lemma 4.3 and Remark 4.1, it suffices to show that the term $I(t)$ grows as $\mathcal{O}(L(t)\sqrt{\log \log L(t)})$ with probability 1. To do so, let $\kappa_p := [I_p]$ denote the quadratic variation of I_p .⁹ Then, the rules of stochastic calculus yield

⁹ Recall here that the quadratic variation of a stochastic process $M(t)$ is the continuous increasing and progressively measurable process, defined as $[M(t)] = \lim_{|\Pi| \rightarrow 0} \sum_{1 \leq j \leq k} (M(t_j) - M(t_{j-1}))^2$, where the limit is taken over all partitions $\Pi = \{t_0 = 0 < t_1 < \dots < t_k = t\}$ of $[0, t]$ with mesh $|\Pi| \equiv \max_j |t_j - t_{j-1}| \rightarrow 0$ [32].

$$\begin{aligned}
d\kappa_p(t) &= dI_p(t) \cdot dI_p(t) \\
&= \lambda^2(t) \sum_{i,j=1}^n \sum_{k=1}^d (X_i(t) - p_i)(X_j(t) - p_j) \sigma_{ik}(X(t), t) \sigma_{jk}(X(t), t) dt \\
&\leq \|X(t) - p\|_2^2 \sigma_*^2 \lambda^2(t) \\
&\leq \text{diam}(\mathcal{X})^2 \sigma_*^2 \lambda^2(t),
\end{aligned}$$

where $\text{diam}(\mathcal{X}) := \max\{\|x' - x\|_2 : x, x' \in \mathcal{X}\}$ denotes the Euclidean diameter of \mathcal{X} . Hence, for all $t \geq 0$, we get the quadratic variation bound

$$\kappa_p(t) \leq \text{diam}(\mathcal{X})^2 \sigma_*^2 \int_0^t \lambda^2(s) ds = \mathcal{O}(L^2(t)). \quad (22)$$

Now, let $\kappa_p(\infty) := \lim_{t \rightarrow \infty} \kappa_p(t) \in [0, \infty]$ and set

$$\tau_p(s) := \begin{cases} \inf\{t \geq 0 : \kappa_p(t) > s\}, & \text{if } s \leq \kappa_p(\infty), \\ \infty, & \text{otherwise.} \end{cases}$$

The process $\tau_p(s)$ is finite, nonnegative, nondecreasing and right-continuous on $[0, \kappa_p(\infty))$; moreover, it is easy to check that $\kappa_p(\tau_p(s)) = s \wedge \kappa_p(\infty)$ and $\tau_p(\kappa_p(t)) = t$ [32, Problem 3.4.5]. Therefore, by the Dambis–Dubins–Schwarz time-change theorem for martingales [32, Theorem. 3.4.6 and Problem. 3.4.7], there exists a standard, one-dimensional Wiener process $(B_p(t))_{t \geq 0}$ adapted to a modified filtration $\tilde{\mathcal{F}}_s = \mathcal{F}_{\tau_p(s)}$ (possibly defined on an extended probability space), and such that $B_p(\kappa_p(t)) = I_p(t)$ for all $t \geq 0$ (except possibly on a \mathbb{P} -null set). Hence, for all $t > 0$, we have

$$\frac{I_p(t)}{S(t)} = \frac{B_p(\kappa_p(t))}{S(t)} = \frac{B_p(\kappa_p(t))}{\sqrt{\kappa_p(t) \log \log \kappa_p(t)}} \times \frac{\sqrt{\kappa_p(t) \log \log \kappa_p(t)}}{S(t)}.$$

By the law of the iterated logarithm [32], the first factor above is bounded almost surely; as for the second, (22) gives $\sqrt{\kappa_p(t) \log \log \kappa_p(t)} = \mathcal{O}(L(t) \sqrt{\log \log L(t)})$. Thus, combining all of the above, we get

$$\frac{I(t)}{S(t)} = \frac{\max_{p \in \mathcal{X}} I_p(t)}{S(t)} = \mathcal{O} \left(\frac{L(t) \sqrt{\log \log L(t)}}{S(t)} \right),$$

so our claim follows from (16).

To complete our proof, note first that the condition $\lim_{t \rightarrow \infty} \eta(t)S(t) = \infty$ implies that $\lim_{t \rightarrow \infty} S(t) = \infty$ (given that $\eta(t)$ is nonincreasing). Thus, by de l'Hôpital's rule and the assumption $\lim_{t \rightarrow \infty} \lambda(t)\eta(t) = 0$, we also get $S(t)^{-1} \int_0^t \lambda^2(s)\eta(s) ds = 0$. Finally, for the last term of (14), consider the following two cases:

1. If $\lim_{t \rightarrow \infty} L(t) < \infty$, we trivially have $\lim_{t \rightarrow \infty} L(t) \sqrt{\log \log L(t)} / S(t) = 0$ as well.

2. Otherwise, if $\lim_{t \rightarrow \infty} L(t) = \infty$, de l’Hôpital’s rule readily yields

$$\lim_{t \rightarrow \infty} \frac{L^2(t)}{S^2(t)} = \lim_{t \rightarrow \infty} \frac{\lambda^2(t)}{2\lambda(t)S(t)} = \frac{1}{2} \lim_{t \rightarrow \infty} \frac{\lambda(t)}{S(t)} = 0,$$

by the boundedness of $\lambda(t)$. Another application of de l’Hôpital’s rule gives

$$\lim_{t \rightarrow \infty} \frac{L^3(t)}{S^2(t)} = \lim_{t \rightarrow \infty} \frac{(L^2(t))^{3/2}}{S^2(t)} = \frac{3}{4} \lim_{t \rightarrow \infty} \frac{\lambda^2(t)L(t)}{\lambda(t)S(t)} = \frac{3}{4} \lim_{t \rightarrow \infty} \frac{\lambda(t)L(t)}{S(t)} = 0,$$

so

$$\limsup_{t \rightarrow \infty} \frac{L(t)\sqrt{\log \log L(t)}}{S(t)} \leq \limsup_{t \rightarrow \infty} \sqrt{\frac{L^3(t)}{S^2(t)}} = 0.$$

The above shows that, under the stated assumptions, the RHS of (14) converges to 0 almost surely, implying in turn that $\bar{X}(t)$ converges to the solution set of $\text{VI}(\mathcal{X}, v)$ with probability 1. □

4.4 Large Deviations

In this section, we study the concentration properties of (SMD) in terms of the dual gap function. As in the previous section, we will assume that (SMD) is issued from the initial condition $(s, y) = (0, 0)$.

First, recall that for every $p \in \mathcal{X}$ we have the upper bound

$$R_p(t) \leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \frac{\sigma_*^2}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds + I_p(t).$$

with $R_p(t)$ and $I_p(t)$ defined as in (19) and (17), respectively. Since $I_p(t)$ is a continuous martingale starting at 0, we have $\mathbb{E}[I_p(t)] = 0$, implying in turn that

$$\mathbb{E}[(v(p), \bar{X}(t) - p)] \leq \frac{\mathcal{D}(h; \mathcal{X})}{S(t)\eta(t)} + \frac{\sigma_*^2}{2\alpha S(t)} \int_0^t \lambda^2(s)\eta(s) \, ds = \frac{K(t)}{2S(t)},$$

where

$$K(t) := \frac{2\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \frac{\sigma_*^2}{\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds. \tag{23}$$

Markov’s inequality therefore implies that

$$\mathbb{P}((v(p), \bar{X}(t) - p) \geq \delta) \leq \frac{1}{\delta} \frac{K(t)}{2S(t)} \text{ for all } \delta > 0. \tag{24}$$

The bound (24) provides a first estimate of the probability of observing a large gap from the solution of (MVI), but because it relies only on Markov’s inequality, it is rather crude. To refine it, we provide below a “large deviations” bound that shows that the ergodic gap process $g(\bar{X}(t))$ is exponentially concentrated around its mean value:

Theorem 4.3 *Suppose (H1)–(H3) hold, and that (SMD) is started from the initial condition $(s, y) = (0, 0)$. Then, for all $\delta > 0$ and all $t > 0$, we have*

$$\mathbb{P} \left(g(\bar{X}(t)) \geq \mathcal{Q}_0(t) + \delta \mathcal{Q}_1(t) \right) \leq \exp(-\delta^2/4), \tag{25}$$

where

$$\mathcal{Q}_0(t) := \frac{K(t)}{S(t)}, \tag{26a}$$

and

$$\mathcal{Q}_1(t) := \frac{\sqrt{\kappa} \sigma_* \text{diam}(\mathcal{X}) L(t)}{S(t)}, \tag{26b}$$

with $\kappa > 0$ a positive constant depending only on \mathcal{X} and $\|\cdot\|$.

The concentration bound (25) can also be formulated as follows:

Corollary 4.2 *With notation and assumptions as in Theorem 4.3, we have*

$$\begin{aligned} g(\bar{X}(t)) &\leq \mathcal{Q}_0(t) + 2\mathcal{Q}_1(t)\sqrt{\log(1/\delta)} \\ &= \mathcal{O}\left(\frac{1}{\eta(t)S(t)}\right) + \mathcal{O}\left(\frac{\int_0^t \lambda^2(s)\eta(s) \, ds}{S(t)}\right) + \mathcal{O}\left(\frac{L(t)\sqrt{\log(1/\delta)}}{S(t)}\right), \end{aligned}$$

with probability at least $1 - \delta$. In particular, if (SMD) is run with parameters $\lambda(t) = (1 + t)^{-a}$ and $\eta(t) = (1 + t)^{-b}$ for some $a, b \in [0, 1]$, we have

$$g(\bar{X}(t)) = \mathcal{O}\left(t^{a+b-1}\right) + \mathcal{O}\left(t^{-a-b}\right) + \mathcal{O}\left(t^{-1/2}\right),$$

with arbitrarily high probability.

To prove Theorem 4.3, define first the auxiliary processes

$$Z(t) := \int_0^t \lambda(s)\sigma(X(s), s) \, dW(s) \quad \text{and} \quad P(t) := Q(\eta(t)Z(t)).$$

We then have:

Lemma 4.4 *For all $p \in \mathcal{X}$ we have*

$$\int_0^t \lambda(s)\langle p - P(s), \sigma(X(s), s) \, dW(s) \rangle \leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \frac{\sigma_*^2}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds. \tag{27}$$

Proof The proof follows the same lines as Lemma 4.3. Specifically, given a reference point $p \in \mathcal{X}$, define the process $\tilde{H}_p(t) := \frac{1}{\eta(t)} F(p, \eta(t)Z(t))$. Then, by the weak Itô formula (33) in Sect. 5, we have

$$\begin{aligned} \tilde{H}_p(t) &\leq \tilde{H}_p(0) \\ &\quad - \int_0^t \frac{\dot{\eta}(s)}{\eta(s)^2} \tilde{H}_p(s) \, ds + \frac{1}{\eta(t)} \int_0^t \langle \xi(s) - p, \dot{\eta}(s)Z(s) \rangle \, ds \\ &\quad + \int_0^t \langle P(s) - p, dZ(s) \rangle + \frac{1}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \|\sigma(X(s), s)\|^2 \, ds \\ &\leq - \int_0^t \frac{\dot{\eta}(s)}{\eta(s)} [h(p) - h(P(s))] \, ds \\ &\quad + \int_0^t \lambda(s) \langle P(s) - p, \sigma(X(s), s) \, dW(s) \rangle + \frac{\sigma_*^2}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds. \end{aligned}$$

We thus get

$$\begin{aligned} \int_0^t \lambda(s) \langle p - P(s), \sigma(X(s), s) \, dW(s) \rangle &\leq \tilde{H}_p(0) + \mathcal{D}(h; \mathcal{X}) \left(\frac{1}{\eta(t)} - \frac{1}{\eta(0)} \right) \\ &\quad + \int_0^t \frac{\lambda^2(s)\eta(s)}{2\alpha} \sigma_*^2 \, ds \\ &\leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \int_0^t \frac{\lambda^2(s)\eta(s)}{2\alpha} \sigma_*^2 \, ds, \end{aligned}$$

as claimed. □

We are now ready to establish our large deviations principle for (SMD):

Proof of Theorem 4.3 For $p \in \mathcal{X}$ and $t > 0$ fixed, we have

$$\begin{aligned} R_p(t) &\leq \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \frac{\sigma_*^2}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds + \int_0^t \lambda(s) \langle p - X(s), \sigma(X(s), s) \, dW(s) \rangle \\ &= \frac{\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \frac{\sigma_*^2}{2\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds + \int_0^t \lambda(s) \langle p - P(s), \sigma(X(s), s) \, dW(s) \rangle \\ &\quad + \int_0^t \lambda(s) \langle P(s) - X(s), \sigma(X(s), s) \, dW(s) \rangle \\ &\leq \frac{2\mathcal{D}(h; \mathcal{X})}{\eta(t)} + \frac{\sigma_*^2}{\alpha} \int_0^t \lambda^2(s)\eta(s) \, ds + \int_0^t \langle P(s) - X(s), \sigma(X(s), s) \, dW(s) \rangle, \end{aligned}$$

where we used (27) to obtain the last inequality. To proceed, let

$$\Delta(t) := \int_0^t \lambda(s) \langle P(s) - X(s), \sigma(X(s), s) \, dW(s) \rangle.$$

The process $\Delta(t)$ is a continuous martingale starting at 0 which is almost surely bounded in L^2 , thus providing an upper bound for $R_p(t)$ which is independent of the reference point $p \in \mathcal{X}$. Indeed, recalling the definition (23) of $K(t)$, we see that

$$R_p(t) \leq K(t) + \Delta(t),$$

so

$$g(\bar{X}(t)) \leq \frac{K(t) + \Delta(t)}{S(t)} \quad \text{for all } t > 0.$$

In turn, this implies that for all $\varepsilon, t > 0$,

$$\mathbb{P}(g(\bar{X}(t)) \geq \varepsilon) \leq \mathbb{P}(\Delta(t) \geq \varepsilon S(t) - K(t)) \quad \text{for all } \varepsilon, t > 0.$$

To prove the theorem, we are left to bound the right-hand side of the above expression. To that end, letting $\rho(t) := [\Delta(t), \Delta(t)]$ denote the quadratic variation of $\Delta(t)$, the Cauchy–Schwarz inequality readily gives

$$\begin{aligned} \mathbb{E}[\exp(\theta \Delta(t))] &= \mathbb{E}[\exp(\theta \Delta(t) - b\rho(t)) \exp(b\rho(t))] \\ &\leq \sqrt{\mathbb{E}[\exp(2\theta \Delta(t) - 2b\rho(t))]} \sqrt{\mathbb{E}[\exp(2b\rho(t))]} \end{aligned}$$

Setting $b = \theta^2$, the expressions inside the first expected value is just the stochastic exponential of the process $2\theta \Delta(t)$. Moreover, a straightforward calculation shows that

$$\begin{aligned} \rho(t) &= \int_0^t \lambda^2(s) \|\sigma(X(s), s) \cdot (P(s) - X(s))\|_2^2 \, ds \\ &\leq \kappa \int_0^t \lambda^2(s) \|\sigma(X(s), s)\|^2 \|P(s) - X(s)\|^2 \, ds \\ &\leq \kappa \sigma_*^2 \text{diam}(\mathcal{X})^2 L^2(t), \end{aligned}$$

where $\kappa > 0$ is a universal constant accounting for the equivalence of the Euclidean norm $\|\cdot\|_2$ and the primal norm $\|\cdot\|$ on \mathcal{X} . The above implies that $\rho(t)$ is bounded over every compact interval, showing that Novikov’s condition is satisfied (see, e.g., [32]). We conclude that the process $\exp(2\theta \Delta(t) - 2\theta^2 \rho(t))$ is a true martingale with expected value 1. Hence, letting $\varphi(t) := \kappa \sigma_*^2 \text{diam}(\mathcal{X})^2 L^2(t)$, we get

$$\mathbb{E}[\exp(\theta \Delta(t))] \leq \sqrt{\mathbb{E}[\exp(2\theta^2 \rho(t))]} \leq \exp(\theta^2 \varphi(t)). \tag{28}$$

Combining all of the above, we see that for all $a > 0$

$$\begin{aligned} \mathbb{P}(\Delta(t) \geq a) &= \mathbb{P}(\exp(\theta \Delta(t)) \geq \exp(\theta a)) \\ &\leq \exp(-a\theta) \mathbb{E}[\exp(\theta \Delta(t))] \quad \text{(Markov inequality)} \\ &= \exp(-a\theta + \theta^2 \varphi(t)) \quad \text{for all } a > 0, \end{aligned}$$

with the last line following from (28). Minimizing the above with respect to θ then gives

$$\mathbb{P}(\Delta(t) \geq a) \leq \exp\left(-\frac{a^2}{4\varphi(t)}\right).$$

Hence, by unrolling Eqs. (26a) and (26b), we finally obtain the bound

$$\begin{aligned} \mathbb{P}(g(\bar{X}(t)) \geq Q_0(t) + \delta Q_1(t)) &\leq \mathbb{P}(\Delta(t) \geq Q_0(t)S(t) + \delta Q_1(t)S(t) - K(t)) \\ &\leq \mathbb{P}\left(\Delta(t) \geq \delta\sqrt{\varphi(t)}\right) \\ &\leq \exp(-\delta^2/4), \end{aligned}$$

for all $\delta > 0$, as claimed. \square

5 Conclusions

This paper examined a continuous-time dynamical system for solving monotone variational inequality problems with random inputs. The key element of our analysis is the identification of a energy-type function, which allows us to prove ergodic convergence of generated trajectories in the deterministic as well as in the stochastic case. Future research should extend the present work in the following dimensions. First, it is not clear yet how the continuous-time method will help us in the derivation of a consistent numerical scheme. A naive Euler discretization might potentially lead to a loss in speed of convergence (see [2]). Second, it is of great interest to relax the monotonicity assumption we made on the involved operator. We are currently investigating these extensions. Third, it is of interest to consider different noise models as well. In particular, it would be interesting to know how the results derived in this paper change when the stochastic perturbation comes from a jump Markov process, or more generally, a Lévy process. This extension would likely need new techniques, and we regard this as an important contribution for future work.

Acknowledgements We thank the co-Editor in Chief, Professor Franco Giannessi, for clarifications on the conceptual differences between Minty and Stampacchia Variational inequalities, and some pointers to the relevant literature. P. Mertikopoulos was partially supported by the French National Research Agency (ANR) Grant ORACLESS (ANR-16-CE33-0004-01) and the COST Action CA16228 “European Network for Game Theory” (GAMENET). The research of M. Staudigl is partially supported by the COST Action CA16228 “European Network for Game Theory” (GAMENET).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix A: Results from Convex Analysis

In this appendix, we collect some simple facts on the analysis of convex differentiable functions with Lipschitz continuous gradients. Denote by $C_L^{1,1}(\mathbb{R}^n)$ the totality of such

functions, with L being the Lipschitz constant of the gradient mapping $\nabla\psi$:

$$\|\psi(y + \delta) - \psi(y)\|_* \leq L\|\delta\|_* \quad \forall y, \delta \in \mathbb{R}^n.$$

Proposition A.1 *Let $\psi \in C_L^{1,1}(\mathbb{R}^n)$ be convex. Then, ψ is almost everywhere twice differentiable with Hessian $\nabla^2\psi$ and*

$$0 \leq \nabla^2\psi(y) \leq L \text{Id}. \quad \text{Leb} - a.e..$$

Proof For every $\psi \in C_L^{1,1}(\mathbb{R}^n)$, the well-known descent lemma ([33], Theorem 2.1.5) implies that

$$\psi(y + \delta) \leq \psi(y) + \langle \nabla\psi(y), \delta \rangle + \frac{L}{2}\|\delta\|_*^2 \quad \forall y, \delta \in \mathbb{R}^n. \tag{29}$$

By Alexandrov’s theorem (see, e.g., [34], Lemma 6.6), it follows that ψ is Leb-almost everywhere twice differentiable. Hence, there exists a measurable set Λ such that $\text{Leb}(\Lambda) = 0$, and for all $\bar{y} \in \mathbb{R}^n \setminus \Lambda$ there exists $(p, P) \in \mathbb{R}^n \times \mathbb{R}_{sym}^{n \times n}$ such that

$$\psi(\bar{y} + y) = \psi(\bar{y}) + \langle p, y \rangle + \frac{1}{2}\langle Py, y \rangle + \theta(\bar{y}, y), \tag{30}$$

where $\lim_{\|y\|_* \rightarrow 0} \frac{\theta(\bar{y}, y)}{\|y\|_*^2} = 0$. We have $p = \nabla\psi(\bar{y})$ and identify P with the a.e. defined Hessian $\nabla^2\psi(\bar{y})$. On the other hand, convexity implies

$$\psi(\bar{y} + y) \geq \psi(\bar{y}) + \langle \nabla\psi(\bar{y}), y \rangle \quad \forall \bar{y}, y \in \mathbb{R}^n. \tag{31}$$

Choosing $y = te$, where $e \in \mathbb{R}^n$ is an arbitrary $\|\cdot\|_*$ -unit vector and $t > 0$, it follows

$$\begin{aligned} -\frac{1}{t^2}\theta(\bar{y}, te) &\stackrel{(31)}{\leq} \frac{1}{t^2}[\psi(\bar{y} + te) - \psi(\bar{y}) - \langle \nabla\psi(\bar{y}), te \rangle] - \frac{1}{t^2}\theta(\bar{y}, te) \\ &\stackrel{(30)}{=} \frac{1}{2}\langle \nabla^2\psi(\bar{y})e, e \rangle \\ &\stackrel{(29)}{\leq} \frac{L}{2} - \frac{1}{t^2}\theta(\bar{y}, te). \end{aligned}$$

Letting $t \rightarrow 0^+$ we get

$$0 \leq \frac{1}{2}\langle \nabla^2\psi(\bar{y})e, e \rangle \leq \frac{L}{2} \quad \forall \bar{y} \in \mathbb{R}^n \setminus \Lambda, \tag{32}$$

which implies $\nabla^2\psi(\bar{y}) \leq L \text{Id}$. □

Appendix B: Results from Stochastic Analysis

The following result is the generalized Itô formula, used in the main text.

Proposition B.1 *Let Y be an Itô process in \mathbb{R}^n of the form*

$$Y_t = Y_0 + \int_0^t F_s \, ds + \int_0^t G_s \, dW(s).$$

Let $\psi \in C_L^{1,1}(\mathbb{R}^n)$ be convex. Then, for all $t \geq 0$, we have

$$\psi(Y_t) \leq \psi(Y_0) + \int_0^t \langle \nabla \psi(Y_s), dY_s \rangle + \frac{L}{2} \int_0^t \|G_s\|^2 \, ds$$

Proof Since $\psi \in C_L^{1,1}(\mathbb{R}^n)$ is convex, Proposition A.1 shows that ψ is almost everywhere twice differentiable with Hessian $\nabla^2 \psi$. Furthermore, this Hessian matrix satisfies $0 \leq \nabla^2 \psi(y) \leq L \text{Id}$, for all $y \in \mathbb{R}^n$ outside a set of Lebesgue measure 0.

Introduce the mollifier

$$\rho(u) := \begin{cases} c \exp\left(\frac{-1}{1-\|u\|_*^2}\right), & \text{if } \|u\|_* < 1, \\ 0, & \text{if } \|u\|_* \geq 1. \end{cases}$$

Choose the constant $c > 0$ so that $\int_{\mathbb{R}^n} \rho(u) \, du = 1$. For every $\varepsilon > 0$ define

$$\begin{aligned} \rho_\varepsilon(u) &:= \varepsilon^{-n} \rho(u/\varepsilon), \\ \psi_\varepsilon(y) &:= \psi \otimes \rho_\varepsilon(y) := \int_{\mathbb{R}^n} \psi(y - u) \rho_\varepsilon(u) \, du. \end{aligned}$$

Then, $\psi_\varepsilon \in C^\infty(\mathbb{R}^n)$ and the standard form of Itô’s formula gives us

$$\begin{aligned} \psi_\varepsilon(Y_t) &= \psi_\varepsilon(Y_s) + \int_s^t \langle \nabla \psi_\varepsilon(Y_r), dY_r \rangle + \frac{1}{2} \int_s^t \text{tr} \left[\nabla^2 \psi_\varepsilon(Y_r) G_r G_r^\top \right] \, dr \\ &= \psi_\varepsilon(Y_s) + \int_s^t \left\langle \int_{\mathbb{R}^n} \nabla \psi(z) \rho_\varepsilon(Y_r - z) \, dz, dY_r \right\rangle \\ &\quad + \frac{1}{2} \int_s^t \int_{\mathbb{R}^n} \text{tr} \left[\nabla^2 \psi(z) G_r G_r^\top \right] \rho_\varepsilon(Y_r - z) \, dr \, dz. \end{aligned}$$

Since $\text{tr}(\nabla^2 \psi(z) G_r G_r^\top) \leq L \|G_r\|^2$, we get

$$\psi_\varepsilon(Y_t) \leq \psi_\varepsilon(Y_s) + \int_s^t \left\langle \int_{\mathbb{R}^n} \nabla \psi(z) \rho_\varepsilon(Y_r - z) \, dz, dY_r \right\rangle + \frac{L}{2} \int_s^t \|G_r\|^2 \, dr.$$

Letting $\varepsilon \downarrow 0$, using the uniform convergence of the involved data, proves the result. □

Applying this result to the dual process of (SMD) and using (32), gives for $F_s = A(s, Y(s, y))$ and $G_s = B(s, Y(s, y))$, the following version of the generalized Itô rule:

$$\psi(Y_t) \leq \psi(Y_0) + \int_0^t \langle \nabla \psi(Y_s), dY_s \rangle ds + \frac{1}{2\alpha} \int_0^t \|B(s, Y(s, y))\|^2 ds \quad (33)$$

References

1. Peypouquet, J., Sorin, S.: Evolution equations for maximal monotone operators: asymptotic analysis in continuous and discrete time. *J Convex Anal.* **17**(3&4), 1113–1163 (2010)
2. Wibisono, A., Wilson, A.C., Jordan, M.I.: A variational perspective on accelerated methods in optimization. *Proc Natl Acad Sci* **113**(47), E7351–E7358 (2016)
3. Nemirovski, A.S., Yudin, D.B.: *Problem Complexity and Method Efficiency in Optimization*. Wiley, New York (1983)
4. Bolte, J., Teboulle, M.: Barrier operators and associated gradient-like dynamical systems for constrained minimization problems. *SIAM J. Control Optim.* **42**(4), 1266–1292 (2003)
5. Alvarez, F., Bolte, J., Brahic, O.: Hessian Riemannian gradient flows in convex programming. *SIAM J. Control Optim.* **43**(2), 477–501 (2004)
6. Attouch, H., Bolte, J., Redont, P., Teboulle, M.: Singular Riemannian barrier methods and gradient-projection dynamical systems for constrained optimization. *Optimization* **53**(5–6), 435–454 (2004)
7. Harker, P.T., Pang, J.S.: Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Math. Program.* **48**(1), 161–220 (1990)
8. Giannessi, F.: On minty variational principle. In: Giannessi, F., Komlósi, S., Rapcsák, T. (eds.) *New Trends in Mathematical Programming, Applied Optimization*, vol. 13, pp. 93–99. Springer, Boston (1998)
9. Facchinei, F., Pang, J.S.: *Finite-Dimensional Variational Inequalities and Complementarity problems*. Springer, Berlin (2003)
10. Ferris, M.C., Pang, J.S.: Engineering and economic applications of complementarity problems. *SIAM Rev.* **39**(4), 669–713 (1997)
11. Scutari, G., Palomar, D.P., Facchinei, F., Pang, J.S.: Convex optimization, game theory, and variational inequality theory. *IEEE Signal Process. Mag.* **27**(3), 35–49 (2010)
12. Rockafellar, R.T., Wets, R.J.B.: *Variational Analysis, A Series of Comprehensive Studies in Mathematics*, vol. 317. Springer, Berlin (1998)
13. Nesterov, Y.: Primal-dual subgradient methods for convex problems. *Math. Program.* **120**(1), 221–259 (2009)
14. Mertikopoulos, P., Zhou, Z.: Learning in games with continuous action sets and unknown payoff functions. *Math. Program.* (forthcoming)
15. Xiao, L.: Dual averaging methods for regularized stochastic learning and online optimization. *J. Mach. Learn. Res.* **11**(Oct), 2543–2596 (2010)
16. Duchi, J.C., Agarwal, A., Wainwright, M.J.: Dual averaging for distributed optimization: convergence analysis and network scaling. *IEEE Trans. Autom. Control* **57**(3), 592–606 (2012)
17. Facchinei, F., Kanzow, C.: Generalized Nash equilibrium problems. *4OR* **5**(3), 173–210 (2007)
18. Scutari, G., Facchinei, F., Palomar, D.P., Pang, J.S.: Convex optimization, game theory, and variational inequality theory in multiuser communication systems. *IEEE Signal Process. Mag.* **27**(3), 35–49 (2010)
19. Cabot, A., Engler, H., Gadat, S.: On the long time behavior of second order differential equations with asymptotically small dissipation. *Trans. Am. Math. Soc.* **361**(11), 5983–6017 (2009)
20. Gadat, S., Panloup, F.: Long time behaviour and stationary regime of memory gradient diffusions. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* **50**(2), 564–601 (2014)
21. Abbas, B., Attouch, H.: Dynamical systems and forward-backward algorithms associated with the sum of a convex subdifferential and a monotone cocoercive operator. *Optimization* **64**(10), 2223–2252 (2015)

22. Krichene, W., Bayen, A., Bartlett, P.: Accelerated mirror descent in continuous and discrete time. In: NIPS '15: Proceedings of the 29th International Conference on Neural Information Processing Systems (2015)
23. Raginsky, M., Bouvrie, J.: Continuous-time stochastic mirror descent on a network: variance reduction, consensus, convergence. In: CDC '13: Proceedings of the 51st IEEE Annual Conference on Decision and Control (2013)
24. Mertikopoulos, P., Staudigl, M.: On the convergence of gradient-like flows with noisy gradient input. *SIAM J. Optim.* **28**(1), 163–197 (2018)
25. Borwein, J.M., Dutta, J.: Maximal monotone inclusions and fitzpatrick functions. *J. Optim. Theory Appl.* **171**(3), 757–784 (2016)
26. Nikaido, H., Isoda, K.: Note on non-cooperative convex games. *Pac. J. Math.* **5**, 807–815 (1955)
27. Bruck Jr., R.E.: On the weak convergence of an ergodic iteration for the solution of variational inequalities for monotone operators in Hilbert space. *J. Math. Anal. Appl.* **61**(1), 159–164 (1977)
28. Nemirovski, A.: Prox-method with rate of convergence $o(1/t)$ for Variational Inequalities with Lipschitz continuous monotone operators and smooth convex-concave Saddle Point Problems. *SIAM J. Optim.* **15**(1), 229–251 (2004)
29. Kwon, J., Mertikopoulos, P.: A continuous-time approach to online optimization. *J. Dyn. Games* **4**(2), 125–148 (2017)
30. Khasminskii, R.Z.: Stochastic Stability of Differential Equations, 2 edn. No. 66 in Stochastic Modelling and Applied Probability. Springer, Berlin (2012)
31. Benaïm, M., Hirsch, M.W.: Asymptotic pseudotrajectories and chain recurrent flows, with applications. *J. Dyn. Differ. Equ.* **8**(1), 141–176 (1996)
32. Karatzas, I., Shreve, S.E.: Brownian Motion and Stochastic Calculus. Springer, Berlin (1998)
33. Nesterov, Y.: Introductory Lectures on Convex Optimization: A Basic Course. No. 87 in Applied Optimization. Kluwer Academic Publishers, Dordrecht (2004)
34. Yong, J., Zhou, X.Y.: Stochastic Controls–Hamiltonian Systems and HJB equations. Springer, Berlin (1999)